Object-Oriented Parallel Programming

Edward Givelberg
Department of Physics and Astronomy
Johns Hopkins University
givelberg@jhu.edu

arXiv:1404.4666v1 [cs.PL] 17 Apr 2014

ABSTRACT

We introduce an object-oriented framework for parallel programming, which is based on the observation that programming objects can be naturally interpreted as processes. A parallel program consists of a collection of persistent processes that communicate by executing remote methods. We discuss code parallelization and process persistence, and explain the main ideas in the context of computations with very large data objects.

Categories and Subject Descriptors

D.1.3 [**Software**]: Concurrent Programming; D.1.4 [**Software**]: Object-oriented Programming

General Terms

Design, Languages

Keywords

Parallel programming, pogramming languages, object-oriented programming languages.

1. INTRODUCTION

Since the early 1990s OpenMP [2] and MPI [1] have been developed as successful frameworks for shared memory and distributed memory programming. Parallel programming has been a subject of extensive research ([5, 6]), which led also to development of many specialized high-level languages, such as the PGAS programming languages [7].

In this paper we introduce an object-oriented framework for parallel programming, which is based on the observation that programming objects can be naturally interpreted as processes. Our research has been motivated by the problem of performing computations with very large data sets. We show, through a sequence of simple examples, that the obejct-oriented framework is sufficiently flexible for performing both local "close to the data" computations and global data-intensive tasks requiring moving data with maximal parallelism.

The increasing gap between the speed of performing arithmetic operations and the speed of moving data has recently led to development of communication-avoiding algorithms [3]. On the other hand, the problem of computing a Fourier transform on a very large (Petascale) three-dimensional array can be considered as a prototype problem where massive and highly parallel data communications are necessary. This problem is motivating the examples of this paper. Since we currently do not have a compiler implementing processes, as they are described in 2, our computation of the Fourier transform imitates this framework using standard C++ and several functions of the MPI 2.0 standard. The results of the large-scale Fourier transform computation will be reported in a separate paper.

We introduce the main ideas informally, demonstrating simple use cases. Our examples are written in C++ pseudocode, but the ideas naturally apply to any object-oriented language. Processes are described in sections 2 and 3. The important topics of parallelization and process persistence are dealt with very briefly in sections 4 and 5, respectively.

2. PROCESSES

Programming objects can be thought of as processes. We introduce this idea using a simple example. Consider a device for storing blocks of unstructured data which are described by the following `Page` class:

```
class Page
{
public:
  Page(int n, unsigned char * data);
  ~Page();
private:
  unsigned char * data;
};
```

A `Page` object stores `n` bytes of data in `data`. The block storage device, defined in the `PageDevice` class, uses a file to store multiple data pages of the same size:

```
class PageDevice
{
public:
  PageDevice(
    string filename,
    int NumberOfPages,
    int PageSize
  );
  ~PageDevice();
  void write(Page * p, int PageIndex);
  void read(Page * p, int PageIndex);
protected:
  string filename;
  int NumberOfPages;
  int PageSize;
private:
  FILE * f;
};
```

The implementation of this class creates a file `filename` of `NumberOfPages * PageSize` bytes. Pages of data are stored in the `PageDevice` object using a `PageIndex` address, where `PageIndex` is between `0` and `NumberOfPages`. The `write` method copies a data page of size `PageSize` to the location with an offset `PageIndex * PageSize` from the beginning of the file `filename`. Similarly, the `read` method reads a page of data stored at a given integer address in the `PageDevice`. A new `PageDevice` object is created as usual:

```
int NumberOfPages = 10;
int PageSize = 1024;  // bytes

PageDevice * PageStore
  = new PageDevice("pagefile",
      NumberOfPages, PageSize);
```

Consider now the situation where multiple computers *machine 0, machine 1, machine 2*, etc. are available and suppose that the following code is executed on *machine 0*.

```
PageDevice * PageStore
  = new(machine 1)
    PageDevice("pagefile", NumberOfPages,
        PageSize);

Page * page = GenerateDataPage();

int PageAddress = 17;
PageStore->write(page, PageAddress);
```

This program creates a `PageDevice` object on the remote computer *machine 1*, generates a page of data and stores it in the `PageDevice` object on *machine 1*.

Superficially, the above program differs from the standard C++ only in the extension of the operator `new`. The new `new` allocates objects on remote machines, using the address of the remote machine specified inside parentheses. This particular choice of syntax is not important and is only used here to illustrate the new idea. No new syntax is needed to execute methods on remote objects.

The construction of a new object on a remote machine creates a new process on that machine. This new process acts as a server which listens on a communications port, accepts commands from the parent process, acting as a client, and sends results back to the client. The client-server protocol is generated by the compiler from the class desccription. Remote pointer dereferencing triggers a sequence of events, that includes several client-server communications, data transfer and execution of code on both the local and the remote machines.

Process semantics extend naturally to simple objects, as shown in the following example:

```
double * data
  = new(machine 2) double[1024];
data[7] = 3.1415;
double x = data[2];
```

When this code is executed on *machine 0*, a new process is created on *machine 2*. This process allocates a block of `1024` doubles and deploys a server that communicates with the parent client running on *machine 0*. The execution of `data[7] = 3.1415;` requires communication between the client and the server, including sending the numbers `7` and `3.1415` from the client to the server. Similarly, the execution of the following command leads to an assignment of the local variable `x` with a copy of the remote double `data[2]` obtained over the network using client-server communications. We emphasize that code execution is sequential: each instruction, and all communications associated with it, is completed before the following instruction is executed.

Access to the `data` block can be provided to several computing processes, leading to an example of a shared memory implementation:

```
const int N = 128;
class ComputingProcess;
ComputingProcess * computer[N];

for (int i = 0;  i < N;  i ++)
  computer[i] = new(machine i)
    ComputingProcess(data);
```

Although the `data` block is shared among the processes, the computation is sequential. In section 4 we show how this computation can be parallelized.

Finally, we remark that the notion of the class destructor in C++ extends natually to process objects: destruction of a remote object causes termination of the remote process and completion of the correspoding client-server communications.

The introduction of processes, accessible by remote pointers, creates an object-oriented framework for parallel programming. Processes exchange information by executing methods on remote objects rather than by passing messages. De-

velopment of communication protocols, assembly and parsing of messages, and much of the associated code optimization, is relegated to the compiler.

3. PROCESS INHERITANCE

Having defined processes as programming objects, it is now straightforward to derive new processes using previously defined processes. We illustrate process inheritance by extending the example of the previous section. Consider a device for storing three-dimensional array blocks of `N1 * N2 * N3` doubles. The `ArrayPage` class below is easily derived from the previously defined `Page` class to handle blocks of structured data.

```
class ArrayPage:
  public Page
{
public:
  ArrayPage(
    int N1, int N2, int N3,
    double * data
  );
  double sum();
private:
  int N1, N2, N3;
}
```

We added the `sum` method in the `ArrayPage` class as an example of a method that uses the array structure of the data. The definition of the derived process `ArrayPageDevice` is straightforward and requires no new syntax.

```
class ArrayPageDevice:
  public PageDevice
{
public:
  ArrayPageDevice(
    string filename,
    int NumberOfPages,
    int n1, int n2, int n3
  ):
    N1(n1), N2(n2), N3(n3),
    PageDevice(
      filename,
      NumberOfPages,
      N1 * N2 * N3 * sizeof(double)
    )
  {}
  double sum(int PageAddress);
private:
  int N1, int N2, int N3;
};
```

Suppose that an `ArrayPageDevice` object has been created on a remote machine using

```
int n1 = 128;
int n2 = 128;
int n3 = 128;
ArrayPageDevice * blocks
  = new(machine 3)
    ArrayPageDevice(
      "array_blocks",
      NumberOfPages,
      n1, n2, n3
    );
```

The sum of all elements of the fourth page can be computed by first copying the entire page to the local machine:

```
int PageAddress = 4;
ArrayPage * page;
blocks->read(page, PageAddress);
double result = page->sum();
```

Alternatively, the sum can be computed on the remote machine and only the result copied to the local machine:

```
double result = blocks->sum(PageAddress);
```

The need to choose between "moving the data to the computation" and "moving the computation to the data" arises often in the context of data-intensive computations. Object-oriented processes provide a simple mechanism for the programmer to make the choice.

4. PARALLEL COMPUTATION

The implied semantics of processes requires the execution of a remote method to complete before continuing with the computation. A large computation may therefore be carried out jointly by several machines, but no computation is carried out in parallel. Nevertheless, parallelism can easily be achieved, as shown in the following example.

A data-intensive application is likely to maintain a large number of devices to store portions of a data set. Such devices may be created using the following code:

```
ArrayPageDevice * device[N];

for (int i = 0;  i < N;  i ++)
  device[i] = new(machine i)
    ArrayPageDevice(
      "array_blocks",
      NumberOfPages,
      n1, n2, n3
    );
```

The program may subsequently request to obtain pages of data for local processing, one from each storage device:

```
ArrayPage * buffer[N];
int page_address[N];
int k[N];
:
for (int i = 0;  i < N;  i ++)
  device[i]->read(
    buffer[k[i]],
    page_address[i]
  );
```

A page is copied from the address `page_address[i]` in the `i`-th device to the `k[i]`-th page in the local `buffer`. The implementation of this code is as follows:

```
for (int i = 0;  i < N;  i ++)
{
  send read command
    to device[i] server on machine i
  send page_address[i]
    to device[i] server on machine i
  receive a page
    from device[i] server on machine i
  copy the received page to buffer[k[i]]
}
```

This loop can be easily parallelized by the compiler, by splitting it into two loops, as follows:

```
for (int i = 0;  i < N;  i ++)
{
  send read command
    to device[i] server on machine i
  send page_address[i]
    to device[i] server on machine i
}

for (int i = 0;  i < N;  i ++)
{
  receive a page from machine i
  copy the received page to buffer[k[i]]
}
```

When each `ArrayPageDevice` in the `device` array is assigned to a different hard drive, the processes in the above example will carry out disk I/O in parallel.

The next example shows that the object-oriented model of parallel programming has rich expressive power. Consider a collection of processes for a joint computation of a Fourier transform.

```
class Array;

class FFT
{
public:
  FFT(int myid): id(myid) { ... }
  void SetGroup(int myN, FFT * myfft)
    { N = myN; fft = myfft; }
  void transform(int sign, Array * a);
  :
  :
private:
  int N;
  int id;
  FFT * fft;
};
```

The master process creates `N` parallel processes and assigns each process a unique id:

```
FFT * fft[N];
for (int id = 0;  id < N;  id ++)
  fft[id] = new(machine id) FFT(id);
```

It informs each process in the group that it is a part of a group of `N` concurrent processes:

```
for (int id = 0;  id < N;  id ++)
  fft[id]->SetGroup(N, fft);
```

Subsequent inter-process communication can be implemented by executing methods on remote objects, as shown in section 2. A parallel FFT computation is carried out as follows:

```
Array * a = CreateDataForTransform();

int sign = -1;     // forward
for (int id = 0;  id < N;  id ++)
  fft[id]->transform(sign, a);
```

In this computation an `Array` object `a` is a complex large data object consisting of multiple processes, exchanging information with the `fft` processes during the computation.

We give a detailed example of an `Array` class in the next section.

Notice that the `myfft` parameter of the `SetGroup` method is a remote pointer to an array of remote processes, so future reference to its members will result in additional communications. The following deep copy implementation of `SetGroup`, which copies the entire remote array of remote pointers to a local array of remote pointers, is preferable:

```
void FFT::SetGroup(int myN, FFT * myfft)
{
  N = myN;
  fft = new FFT * [N];
  for (int i = 0;  i < N;  i ++)
    fft[i] = myfft[i];  // remote copy
}
```

Encapsulation, which is an important feature of object-oriented programming, clarifies relationships between objects, facilitating parallelization of method execution across distinct objects. Whenever possible, programs should be automatically parallelized by the compiler, without the use of OpenMP-style directives. Such parallelization may expose subtle programming bugs, but in object-oriented programs these should be corrected by clarifying the objects' interfaces.

In this paper we only consider a few examples, which are all trivially parallelizable. Parallel processes are naturally synchronized at the end of the `for` loop, however since these processes may be accessing common objects, an explicit compiler-supported barrier method for arrays of objects may be useful. For example, the processes belonging to the `fft` array can be synchronized with `fft->barrier();`

5. PERSISTENT PROCESSES

In this section we develop an example of the `Array` class, which was introduced above. The `Array` class provides methods for computation with an array object that requires a large number of hardware devices for its storage. A typical example would be a half-petabyte-sized array, stored on hundreds of hard-drives that are attached to multiple computing nodes, which are interconnected by a fast network. For the code in our examples to be valid in the context of large data objects, many instances of the `int` type should be replaced by `size_t`. Nevertheless, for simplicity we continue to use only `int` below.

The `Array` class implements a three-dimensional array of `double` numbers, indexed on the domain
`[0...N1 - 1] * [0...N2 - 1] * [0...N3 - 1]`.
Our storage method is to break up the domain into rectangular blocks of size `n1 * n2 * n3`, using an `ArrayPage` object for each block.

```
typedef
  vector<ArrayPageDevice *>
  BlockStorage;
```

A `BlockStorage` object represents the available hardware storage, where array data pages are stored. A `PageMap` maps logical array page addresses to physical addresses within a `BlockStorage` object:

```
typedef
  struct { int device_id; int index; }
  PageAddress;

struct PageMap
{
  virtual PageAddress
    PhysicalPageAddress(
      int i1, int i2, int i3
    ) const = 0;
};
```

The device_id of a physical page address identifies the ArrayPageDevice in the BlockStorage object and the index variable determines the page address within the ArrayPageDevice. Each ArrayPageDevice process of the BlockStorage object should be assigned to a different hard disk. The PageMap describes the array data layout and is crucial in determining the I/O patterns of the computation. The following class describing array subdomains will be useful in the definition of the Array class.

```
class Domain
{
public:
  Domain(
    int N11, int N12,
    int N21, int N22,
    int N31, int N32
  );
:
};
```

In the definition of the Array class below we provide, in addition to the constructor, a read and a write methods to access a portion of the array defined by the specified domain object. The Array class is a client process for performing computations on a small subdomain of the array data. An application may deploy multiple coordinating Array client processes in parallel.

```
class Array
{
public:
  Array(
    int N1, int N2, int N3,
    int n1, int n2, int n3,
    BlockStorage data,
    PageMap map
  );

  void read(
    double * subarray,
    Domain * domain
  );
  void write(
    double * subarray,
    Domain * domain
  );

  double sum(Domain * domain);

private:
  int N1, N2, N3;   // array sizes
  int n1, n2, n3;   // page sizes
  BlockStorage data;
  PageMap map;
};
```

At runtime the read method assembles the data in subarray using multiple reads of ArrayPage objects from data. The subarray array should be small enough to fit within the memory of the processor. Similarly, the write method updates the corresponding ArrayPage objects in data. The choice of the PageMap determines the degree of parallelism of these I/O operations.

The sum method is an example of an array computation. Its implementation uses the ArrayPageDevice::sum method for every ArrayPage object corresponding to the domain. The partial sums are computed by the data server processes and combined together by the Array client. The sum of the elements of the entire array can be computed by using the Array client in a loop over array subdomains, and by deploying multiple Array clients in parallel. As mentioned before, the PageMap determines the degree of parallelism of the computation, amd its construction should take into account that in a large scale array computation multiple Array processes will run in parallel, communicating with the collection of processes of the data object.

The Array example demonstrates a method for constructing large data sets and performing computations with them. To complete the picture, applications must be able to access previously constructed data sets. In our view large data objects are described as collections of *persistent processes*. Persistent processes are objects that can be destroyed only by explicitly calling the destructor. The runtime system is responsible for storing process representation, and activating and de-activating processes, as needed. Processes can be accessed using a symbolic object address, similar to addresses used by the Data Access Protocol (DAP) [4], for example:

```
PageDevice * page_device =
  "http://data/set/PageDevice/34";
```

Additional research is needed for implementation of process persistence. Here we note that the combination of inheritance and persistence leads to interesting use cases, such as the following example: We add the constructor

```
ArrayPageDevice(PageDevice * page_device);
```

to the ArrayPageDevice class, so that a new process with a pointer to an existing process can be created as follows:

```
ArrayPageDevice * new_device =
  new ArrayPageDevice(page_device);
```

The new_device process may co-exist and communicate with the page_device process, or it may use a copy constructor to copy the state of page_device and subsequently shut it down using:

```
delete page_device;
```

6. CONCLUSION

In this paper we have shown that programming objects have a natural interpretation as processes, and have described the resulting object-oriented framework for parallel programming. In our view a parallel program consists of a collection of persistent processes, which, in general, represent different

programming objects. The processes communicate by executing methods on remote objects. The resulting framework is rich enough to include shared memory and distributed memory programming, as well as other programming models (client-server applications, map-reduce, etc.).

Processes can be added to any object-oriented language, and should be useful in computations with large data sets, operating system design and scientific applications. In our opinion the process-oriented programming style facilitates creating automatically parallelizable code. In this paper we touched only briefly on the important subjects of code parallelization and process persistence. These topics, as well as issues of implementation and optimization require further research.

7. ACKNOWLEDGEMENT

This research was partially funded by a grant from Intel Corp.

8. REFERENCES

[1] MPI: A Message-Passing Interface Standard, version 3.0. `http://www.mpi-forum.org/`.

[2] OpenMP 4.0 Specifications. `http://openmp.org/wp/openmp-specifications/`.

[3] J. Demmel, M. Hoemmen, M. Mohiyuddin, and K. Yelick. Avoiding communication in sparse matrix computations. In *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, pages 1–12. IEEE, 2008.

[4] J. Gallagher, N. Potter, T. Sgouros, S. Hankin, and G. Flierl. The data access protocol-DAP 2.0, 2004.

[5] M. Herlihy and N. Shavit. *The Art of Multiprocessor Programming, Revised Reprint*. Elsevier, 2012.

[6] T. G. Mattson, B. A. Sanders, and B. L. Massingill. *Patterns for parallel programming*. Pearson Education, 2004.

[7] K. Yelick, D. Bonachea, W.-Y. Chen, P. Colella, K. Datta, J. Duell, S. L. Graham, P. Hargrove, P. Hilfinger, P. Husbands, et al. Productivity and performance using partitioned global address space languages. In *Proceedings of the 2007 international workshop on Parallel symbolic computation*, pages 24–32. ACM, 2007.